



(11) **EP 1 087 575 A1**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
28.03.2001 Bulletin 2001/13

(51) Int Cl.7: **H04L 12/46**

(21) Application number: **99307550.6**

(22) Date of filing: **24.09.1999**

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
 MC NL PT SE**
 Designated Extension States:
AL LT LV MK RO SI

(71) Applicant: **BRITISH TELECOMMUNICATIONS
 public limited company**
London EC1A 7AJ (GB)

(72) Inventor: **The designation of the inventor has not
 yet been filed**

(74) Representative:
Semos, Robert Ernest Vickers et al
BT Group Legal Services,
Intellectual Property Department,
8th Floor, Holborn Centre
120 Holborn
London EC1N 2TE (GB)

(54) **Packet network interfacing**

(57) A method of establishing a tunnel across an IPv4 domain for the transport of packets from a source host on one IPv6 domain to a destination host on another IPv6 domain, there being respective interfaces between the IPv4 domain and the IPv6 domains. The source host sends a normal IPv6 address request to its local DNS server, which relays it to an IPv6 name server in the other IPv6 domain. The response message, containing the true IPv6 address of the destination is received at the remote interface, which appends to the resulting protocol converted DNS response message a first additional record containing the true IPv6 address, and a second additional record containing the IPv4 address of that interface. Upon receipt at the near interface, the additional records are stripped off, their contents stored in an entry of a table, and the true IPv6 address written into the address record of the resulting IPv6 DNS response message. When the near interface receives a packet from an IPv6 host, it checks whether the destination address matches an entry of its table, and if so sends the packet to the encapsulator together with the IPv4 address of the remote interface. The remote interface extracts the source address and the address of the encapsulating interface and stores these in an entry in its corresponding table for use in encapsulating return packets to the source. If, however, the destination address is recognised as being of IPv4-compatible or IPv4-mapped format, the packet is sent to a protocol converter.




Fig. 2

EP 1 087 575 A1

Description

[0001] This invention relates to the establishment of a tunnel entrance for packets crossing the border between a first network operating in accordance with a first transmission protocol and having network addresses in accordance with a first addressing convention, herein referred to as first type addresses, and a second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention, herein referred to as second type addresses, and relates particularly, but not exclusively, to communication between hosts in respective Internet Protocol version 6 (IPv6) domains separated by an Internet Protocol version 4 (IPv4) domain.

[0002] Herein, the terms packet and message are used interchangeably and have the same meaning, and the term Internet domain is used as a specific example of a network.

[0003] In Internet technology, it had become apparent that the initial transport protocol, IPv4, needed to be enhanced, principally to increase the available address space and add a hierarchical address structure. The result was IPv6 which has a simplified header format compared with IPv4, but which uses 128 bit addresses as compared with 32 bit addresses used in IPv4.

[0004] Readers wishing to have an overview of this general area might like to access a list of Internet-Drafts, which are working documents of the IETF, at <http://www.ietf.org/1id-abstracts.txt>, and a particularly relevant document is "A Guide to the Introduction of IPv6 in the IPv4 World" <draft-ietf-ngtrans-introduction-to-ipv6-transition-01.txt>, also referred to as "Guide to IPv6 Transition".

[0005] As mentioned, the present invention relates to tunnelling. Known tunnelling techniques are of two types: configured and automatic.

[0006] A configured tunnel is created by manual configuration of a tunnelling interface between an IPv6 domain and an IPv4 domain such that all packets received from the IPv6 domain are encapsulated in IPv4 packets addressed to a specific tunnel end point, i.e. the tunnelling interface between the IPv4 domain and the remote IPv6 domain containing the destination IPv6 host.

[0007] An automatic tunnel on the other hand does not need manual configuration: the tunnel end points are automatically determined. Several automatic tunnelling mechanisms are currently in development within the IETF, these being known in the art as, 6over4, 6to4, Dynamic Tunnelling, and Tunnel Broker.

[0008] For more detailed information on 6over4 the reader can obtain from the IETF, the document known as RFC2829, or any variant thereof, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", by B. Carpenter and C. Jung, March 1999.

[0009] For more detailed information on 6to4 the reader can obtain from the IETF, the document known

as draft-ietf-ngtrans-6to4-02.txt, or any variant thereof, "Connection of IPv6 Domains via IPv4 Clouds without Explicit Tunnels", by B. Carpenter and K. Moore.

[0010] For more detailed information on Dynamic Tunnelling the reader can obtain from the IETF, the document known as draft-ietf-ngtrans-dti-00.txt.

[0011] For more detailed information on Tunnel Broker the reader can obtain from the IETF, the document known as draft-ietf-ngtrans-broker-00.txt.

[0012] These known automatic tunnelling mechanisms use a variety of techniques to enable the tunnel to be automatically established:

- 6over4 Multicast
- 6to4 Special IPv6 address in which the top level aggregator (TLA) contains an identifier for the 6to4 tunnelling mechanism, and the next level aggregator (NLA) contains the IPv4 address of the tunnel end point
- Dynamic Tunnelling via DNS
- Tunnel Broker www based tool

[0013] In accordance with a first aspect of the present invention, there is provided a method of establishing a tunnel from a first interface between a first network and a second network to a second interface between the second network and a third network, the first and third networks operating in accordance with a first transmission protocol and having addresses in accordance with a first addressing convention, and the second network operating in accordance with a second transmission protocol and having addresses in accordance with a second addressing convention, the tunnel being for the transport of messages from a first host on the first network to a second host on the third network, the method comprising the steps of:

sending from the first host an address request message in accordance with the first transmission protocol, referred to herein as a first type address request message, containing the name of the second host;

upon receipt of the address request message at a name to address conversion system of the third network, returning an address response message in accordance with the first transmission protocol, referred to herein as a first type address response message, and containing the address of the second host in a response address field;

upon receipt of that first type address response message at the second interface,

converting it to an address response message in accordance with the second transmission protocol, referred to herein as a second type address response message, and augmenting that converted second type address response message by fields respectively

containing the address of the second interface in accordance with the second addressing convention and the address of the second host in accordance with the first addressing convention; and

upon receipt of that augmented that converted second type address response message at the first interface,

converting it to a first type address response message,
retrieving the contents of the augmenting fields,
storing at the first interface a mapping of the retrieved address of the second host and the retrieved address of the second interface for use in encapsulating messages from the first host addressed to the second host, and
replacing the content of the response address field of the resulting first type address response message by the retrieved address of the second host.

[0014] Preferably, for establishing a tunnel in the reverse direction for the transport of messages from the second host to the first host, the method comprises the further steps of:

upon receipt at the first interface of a message from the first host addressed to the second host, encapsulating that received message in accordance with the first mapping; and
upon receipt of the encapsulated message at the second interface,

un-encapsulating that received encapsulated message,
retrieving from the encapsulating header the address of the first interface in accordance with the second addressing convention,
retrieving from the un-encapsulated message the address of the first host in accordance with the first addressing convention, and
storing at the second interface a mapping of the retrieved address of the first host and the retrieved address of the first interface for use in encapsulating messages from the second host addressed to the first host.

[0015] There may be included the steps of setting a time to live for a said stored mapping, and rendering that stored mapping unuseable upon the expiry of the time to live, preferably deleting that stored mapping.

[0016] In accordance with a second aspect of the present invention, there is provided a method of sending packets from a first host on a first network via a second network to a second host on a third network, the first and third networks operating in accordance with a first

transmission protocol and having addresses in accordance with a first addressing convention, and the second network operating in accordance with a second transmission protocol and having addresses in accordance with a second addressing convention, comprising the steps of:

establishing a tunnel from a first interface between the first network and the second network to a second interface between the second network and the third network in accordance with the method of the first aspect of the present invention;

upon receipt by the first host of the resulting first type address response message from the first interface, retrieving the content of the response address field;

generating one or more packets for transmission having a header including source and destination address fields, the source address field containing the address of the first host, and the destination address field containing the retrieved content of the response address field;

sending the or each generated packet to the first interface;

for the or each said generated packet received by the first interface, accessing the stored mappings in accordance with the destination address of the received generated packet,

retrieving the stored interface address of the mapping whose retrieved host address matches the destination address of the received generated packet,

generating an encapsulated packet having a payload formed by the received generated packet, and having a header including source and destination address fields, the source address field containing the address of the first interface, and the destination address field containing the retrieved interface network address, sending the encapsulated packet to the second interface; and

for the or each encapsulated packet received by the second interface, un-encapsulating the received encapsulated packet to recover the original generated packet forming its payload, and

sending that recovered packet to the second host.

[0017] There may be included the step of storing the retrieved content in association with the name of the second host.

[0018] There may be included the steps of setting a time to live for the stored retrieved content, and rendering that stored retrieved content unuseable upon the expiry of the time to live, preferably deleting the stored retrieved content.

[0019] Protocol converters are known for enabling an IPv6 host to send messages to an IPv4 host. When a new IPv6 host is activated in an IPv6 domain, it employs the technique known as Neighbourhood Discovery (ND) to find out the identity of hosts with whom it can communicate directly. It broadcasts an ND message containing its IPv6 network address, and each host that receives it sends a reply message containing that host's IPv6 network address. Since the domain uses an underlying transport mechanism, say Ethernet using media access control (MAC) addresses, each host receiving the ND message will retrieve the new host's IPv6 network address and also the MAC address of the new host, and the new host will retrieve from each reply message the sending host's IPv6 network address and its MAC address.

[0020] The new host now constructs an ND table in which each entry corresponds with a neighbouring host and comprises a first part in the form of the 128 bit IPv6 address of that neighbouring host, and a second part in the form of the associated MAC address.

[0021] The interface device (containing the protocol converter) between that IPv6 domain and an adjacent IPv4 domain will also have received the ND message and sent a reply message, and the new host will have made a special Default entry in the ND table having its first part formed by 128 zeros (in variants, these are all ones) and its second part formed by the MAC address of that interface device.

[0022] Thus, when that new host wants to send a message to one of the other hosts in its domain, it constructs an IPv6 message and accesses its ND table to retrieve the MAC address associated with the destination address. The message is then encapsulated within an Ethernet packet in known manner and sent via the underlying Ethernet transport mechanism to the destination host.

[0023] If, on the other hand, the host constructs an IPv6 message having its destination address in the form of an IPv4-compatible or IPv4-mapped address, i.e. a message intended for an IPv4 host in the adjacent IPv4 domain, this destination address will not be found in the ND table. In this situation, the accessing algorithm will return the MAC address of the Default entry, and the message will be sent to the protocol converter.

[0024] Protocol converters can only convert (or translate) between corresponding fields of the headers of the IPv6 and IPv4 messages. Where a field in the header of, say, an IPv6 message has no corresponding field in the header of an IPv4 message, or vice versa, the information in that field will be lost in the protocol conversion process.

[0025] As mentioned above, tunnelling techniques are known for enabling IPv6 hosts to communicate amongst themselves when they are in spaced-apart domains. In this case, the interface device contains a tunnelling mechanism instead of a protocol converter. It will be appreciated that before now, for an IPv6 host to be

able to communicate both with IPv4 hosts and with remote IPv6 hosts, it was necessary for that IPv6 host and the domain interface to be dual stack, i.e. having both IPv4 and IPv6 communication capability. If an IPv6 host was not dual stack, its accessing algorithm would return only a single MAC address for the Default entry. This would be the MAC address of the input port of a protocol converter, if the network administration had decided that the IPv6 host is to be able to communicate with IPv4 hosts, or it would be the MAC address of the input port of a tunnelling mechanism, if the network administration had decided that the IPv6 host is to be able to communicate with IPv6 hosts. That Default entry MAC address could not be a common input address for both the protocol converter and the tunnelling mechanism.

[0026] Methods of tunnel establishment in accordance with the present invention are advantageous in that the source host need only send a standard address request message in its own protocol, e.g. if the first network transmission protocol is IPv6 then that message will be an IPv6 DNS address request message, and the establishment of the tunnel across the intervening IPv4 domain is automatic and completely autonomous, requiring nothing special from the source host. Thus, a company that has an existing IPv4 domain can make the transition from full IPv4 implementation to full IPv6 implementation by introducing isolated IPv6 domains using standard IPv6 technology.

[0027] Specific embodiments of the present invention will now be described with reference to the drawings in which:

Figure 1 is a schematic diagram of an IPv4 domain interfaced with two isolated IPv6 domains;

Figure 2 is a schematic diagram of a border router; Figure 3 is a schematic diagram of an IPv6 DNS Response message;

Figure 4 is a schematic diagram of an IPv4 DNS Response message resulting from conversion of the IPv6 DNS Response message of Figure 3;

Figure 5 is a schematic diagram of an IPv6 DNS Response message resulting from conversion of the IPv4 DNS Response message of Figure 4; and

Figure 6 is a schematic diagram showing the format of special IPv6 addresses used with the 6to4 tunnelling technique.

[0028] In Figure 1, an IPv4 domain 10 separates a first IPv6 domain 12, constituting in accordance with the present invention a first network operating in accordance with a first transmission protocol and having network addresses in accordance with a first addressing convention, from a second IPv6 domain 14, constituting in accordance with the present invention a third network also operating in accordance with the first transmission protocol and having network addresses in accordance with the first addressing convention. Hosts in the IPv4 domain 10 are IPv4 only, and hosts in the IPv6 domains

12 and 14 are IPv6 only. The IPv4 domain 10 constitutes in accordance with the present invention a second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention. For simplifying the drawings, no IPv4 hosts are shown, and in each of the IPv6 domains 12 and 14 only one IPv6 host (28 and 30, respectively, as referred to later) is shown.

[0029] The first IPv6 domain 12 is connected to the IPv4 domain 10 via a border router 16A, also referred to as an ingress interface with respect to the first IPv6 domain 12 and constituting an interface of the present invention. The second IPv6 domain 14 is connected to the IPv4 domain 10 via a border router 16B, also referred to as an egress interface with respect to the first IPv6 domain 12. The border router 16B is identical to the border router 16A.

[0030] The IPv4 domain 10 contains a complete domain name system (DNS) 20 including a plurality of DNS servers 22, of which only two DNS servers 22A and 22B are shown, and the IPv6 domains 12 and 14 contain respective DNS servers 24 and 26.

[0031] Suppose that a host 28 in the first IPv6 domain 12 wishes to send a packet to a host 30 in the second IPv6 domain 14. Thus, for this transaction, the host 28 is referred to as the source host 28 and the host 30 is referred to as the destination host 30.

[0032] The source host 28 knows the name of the destination host 30, so it constructs in known manner an IPv6 DNS Request message (not shown) requesting the IPv6 address of the destination host 30. The source host 28 sends this DNS Request message as a recursive request to its local DNS server, which in this embodiment is the DNS server 24. The DNS server 24 will, in known manner, send a number of iterative DNS Request messages (not shown) to the DNS 20 until it learns about the DNS server 26. Finally, a DNS Request message (not shown) will go to the DNS server 26 requesting the IPv6 address of the destination host 30. As the DNS request passes from the first IPv6 domain 12 through the border router 16A to the IPv4 domain 10, it is processed by a protocol converter (PC) 32A (see Figure 2) and undergoes IPv6/IPv4 translation. Correspondingly, as the DNS request passes from the IPv4 domain 10 through the border router 16B to the second IPv6 domain 14, it is processed by a protocol converter 32B and undergoes IPv4/IPv6 translation.

[0033] The protocol converters 32A and 32B conform to a specification known as Network Address Translation-Protocol Translation (NAT-PT). They translate between IPv4 and IPv6 addresses and keep state during the time of the session, and translation between IPv4 and IPv6 DNS Request messages and DNS Response messages, including translation of IP headers and DNS payload information is controlled by an application layer gateway (ALG). In the art, an alternative term for a DNS Response message is a DNS Answer message.

[0034] For more detailed information the reader can

obtain from the Internet Engineering Task Force (IETF), the document known as draft-ietf-ngtrans-natpt-05.txt, or any variant thereof, "Network Address Translation - Protocol Translation (NATPT)", by G. Tsirtsis and P. Srishuresh.

[0035] The DNS server 26 responds to the DNS Request message for the IPv6 address of the destination host 30 with a IPv6 DNS Response message 34 (see Figure 3) having a conventional format of destination address field 36, source address field 38, and response address record 40 that contains the IPv6 address of the destination host 30.

This IPv6 DNS Response message 34 travels through the second IPv6 domain 14 to the border router 16B where it becomes converted to an IPv4 DNS Response message 42 (see Figure 4) comprising destination address field 44, source address field 46, response address record 48 and, in accordance with the present invention, additional records 50 and 52, and this message 42 travels through the IPv4 domain 10 to the border router 16A where it becomes converted to an IPv6 DNS Response message 54 comprising destination address field 56, source address field 58 and response address record 60, and sent to the source host 28. The route that the DNS Response message (in its various forms, i.e. 34, 42 and 54) takes in the domains 10, 12 and 14 depends on the DNS configuration in each of the domains, but it will have to pass through the border router 16B and the border router 16A in that order.

[0036] For convenience, the terms "field" and "record" are used synonymously and interchangeably in this description, although in the art a field is generally taken to be a component part of a record.

[0037] When the IPv6 DNS Response message 34 is received by the border router 16B via its IPv6 network interface controller 62B, it is fed in parallel to the protocol converter 32B and to a controller 64B, and also to an encapsulator 86B and a 6to4 encapsulator 90B. The controller 64B is connected via a control line 65A to control inputs of the protocol converter 32B, the encapsulator 86B and the 6to4 encapsulator 90B and by placing a suitable address on the control line 65A selects the appropriate one of these devices.

[0038] The controller 64B (i) recognises that the received message is a DNS Response message and enables the protocol converter 32B, (ii) retrieves the IPv6 address from the response record 40 and writes this message to a storage location 66B formed by part of its internal memory 68B; (iii) writes the IPv4 address of the border router 16B, i.e. the IPv4 address of the tunnel terminating endpoint on the border router 16B, to a storage location 70B also formed by part of its internal memory 68B; (iv) receives from the protocol converter 32B the converted DNS Response message 42; appends the content of the storage location 70B as the first additional record 50, and the content of the storage location 68B as the second additional record 52; and (v) sends the resulting IPv4 DNS Response message 42 to an

IPv4 network interface controller 72B for transmission over the IPv4 domain 10 to the border router 16A.

[0039] In a variant, the received IPv6 DNS Response message 34 is fed only to the controller 64B, which writes this message to a storage location 74B formed by part of its internal memory 68B. The controller 64B then (i) retrieves the IPv6 address from the response record 40 and writes this message to the storage location 66B; (ii) writes the IPv4 address of the border router 16B to the storage location 70B; (iii) generates a modified IPv6 DNS Response message by retrieving the content of the storage location 66B and appending the content of the storage location 70B as the first additional record 50, and the content of the storage location 66B as the second additional record 52; and (iv) sends this modified IPv6 DNS Response message to the protocol converter 32B. The ALG of the protocol converter 32 processes only the header and address response record of the DNS Response message to produce the resulting IPv4 DNS Response message 42, i.e. it allows the additional records to remain unchanged.

[0040] It will be appreciated that the feeding of the received message directly to the protocol converter 32B, and the sending of an enabling signal on the control line 65A to the protocol converter 32B by the controller 64B is logically equivalent to the sending of the received message to the protocol converter 32B by the controller 64B in the above variant, and constitutes sending the message to the protocol converter 32B in accordance with the present invention.

[0041] When the IPv4 DNS Response message 42 is received by the border router 16A via its IPv4 network interface controller 72A, it is fed in parallel to the protocol converter 32A and to a controller 64A.

[0042] The controller 64A (i) receives from the protocol converter 32A the output IPv6 DNS Response message comprising destination address field 56, source address field 58, response address record 60, and additional records 50, 52; and (ii) retrieves the IPv6 address from the second additional record 52, i.e. the true IPv6 address of the destination host 30, and inserts it into the response address record 60 of that output message instead of the IPv4-compatible IPv6 address for the destination host 30, which the protocol converter 32A had generated. The controller 64A then strips off the additional records 50, 52, and sends the resulting IPv6 DNS Response message 54 (see Figure 5) to an IPv6 network interface controller 62A of the border router 16A for transmission to the source host 28.

[0043] Additionally, the controller 64A is arranged to retrieve the IPv4 address of the tunnel terminating endpoint from the first additional record 50, to create a mapping of the IPv6 address of the destination host 30 to the address of the IPv4 tunnel terminating endpoint, and to store this mapping, i.e. create an entry, in an IPv6/tunnel endpoint table 76A formed by part of an internal memory 68A of the controller 64A. This table is referred to as a look-up table.

[0044] In a variant, the additional records 50, 52 are stripped from the DNS Response message prior to the retrieval of their contents. In another variant, the additional records remain attached to the DNS Response message, but this is not as efficient as stripping the additional records. In this present embodiment, each entry in the IPv6/tunnel endpoint table 76A comprises a first element 78A comprising the IPv6 address of a corresponding destination host 30, a second element 80A comprising the IPv4 address of the tunnel terminating endpoint, i.e. of the border router 16B, and third and fourth elements 82A and 84A, to be described later.

[0045] Upon receipt of the resultant IPv6 DNS Response message 54, the source host 28 retrieves the IPv6 address from its address record 60 and stores it in its internal memory for use in sending data packets to the destination host 30.

[0046] In known manner, the source host 28 generates for each of these data packets a header including a source address field and a destination address field, and writes the retrieved IPv6 address into the destination address field.

[0047] Upon receipt of each of these data packets at the border router 16A, the controller 64A retrieves the destination address, and, in accordance with the retrieved destination address, accesses the IPv6/tunnel endpoint table 76A. If there is a match with the contents of a first element 78A of an entry, the controller 64A retrieves the corresponding IPv4 tunnel terminating endpoint from the second element 80A of that entry, and commands an encapsulator 86A to encapsulate the packet in an IPv4 packet addressed to the border router 16B using the retrieved IPv4 tunnel terminating endpoint address. In this embodiment, the encapsulator 86A is arranged to receive the packet directly from the IPv6 network interface controller 62A of the border router 16A, but it does not perform encapsulation until commanded by the controller 64A. In a variant, the controller 64A receives the packet directly from the IPv6 network interface controller 62A and passes it to the encapsulator 86A if there is a match. In practice, when the border router 16A receives a packet the controller 64A writes it into a storage location of its internal memory 68A, and the controller 64A will pass to the encapsulator 86A the address of the relevant storage location, together with an instruction for the encapsulator 86A to access that storage location.

[0048] Upon receipt of the encapsulating IPv4 packet at the border router 16B, an un-encapsulator 88B of the border router 16B strips off the IPv4 header and retrieves the payload of that IPv4 packet, i.e. un-encapsulates the original IPv6 packet from the source host 28, and sends that IPv6 packet to the destination host 30. The controller 64B also creates a mapping (in its IPv6/tunnel endpoint table 76B) of the IPv6 address of the source host 28 and the tunnel originating endpoint, i.e. the IPv4 address of the originating border router 16A, these being respectively retrieved from the source

address field of the IPv6 header and the source address field of the IPv4 header.

[0049] When the destination host 30 returns a Reply packet, a controller 64B of the border router 16B retrieves the destination address, "IPv6 host 28", from the received Reply packet, accesses its IPv6/tunnel endpoint table 76B in accordance with the retrieved destination address (i.e. seeking a matching element 78B), retrieves the corresponding IPv4 address (element 80B), and commands an encapsulator 86B to encapsulate the Reply packet in an IPv4 packet addressed to an un-encapsulator 88A of border router 16A using the IPv4 tunnel originating endpoint address just retrieved from the element 80B of the IPv6/tunnel endpoint table 76B. Upon receipt of this IPv4 packet at the border router 16A, the un-encapsulator 88A performs un-encapsulation to retrieve the Reply packet, and the border router 16A then sends the retrieved Reply packet to the source host 28.

[0050] The source host 28 and the destination host 30 are now in a session in which IPv6 packets pass between them via the tunnel just established between the border routers 16A and 16B.

[0051] The above described mechanism provides for an IPv6 host, which is in an isolated IPv6 domain, to communicate with another IPv6 host, which is in another isolated IPv6 domain, via an intermediate IPv4 domain, without any knowledge of where that other IPv6 host is, and without the source IPv6 host needing to do anything different from a standard communication procedure with another IPv6 host within its own IPv6 domain. The DNS server local to the source IPv6 host makes a Request via the IPv4 domain to the IPv6 DNS server that is on the same network as the destination IPv6 host, and the border routers automatically set up respective mappings associating the tunnel endpoint and the IPv6 address of the IPv6 hosts behind the border routers.

[0052] In this embodiment, the host 28 stores the address retrieved from the DNS Response for a short time. The host 28 records the time at which the address is stored and deletes that stored address twenty four hours later. Alternatively, other methods can be employed, e.g. the address can be deleted at midnight, when the day changes, or the address can be stored with a "time to live" value which is the start value of a count down timer. Similarly, the border router 16A deletes the entry in its IPv6/tunnel endpoint table 76A after a short time. By preventing permanent storage of the address of the host 30, this avoids the risk that the tunnel terminating endpoint becomes out of date for transmissions to the host 30, e.g. the border router 16B is out of service, and that the border router 16A continues to send encapsulated packets to that out of date address. In a variant, the address is deleted after transmission of the last packet of the session. The intent of all the above alternatives is, to a greater or lesser extent, to ensure that tunnels across the IPV4 domain 10 are always up to date.

[0053] In alternative embodiments, some of the en-

tries in the IPv6/tunnel endpoint table 76A can be created by the administrative personnel of the network operator. This is known as manual configuration of a tunnel, and the tunnel is permanent until changed at a later date by the administrative personnel.

[0054] As shown in Figure 2, the border router 16A also comprises a 6to4 tunnelling encapsulator 90A (and 6to4 tunnelling un-encapsulator 92A) and can thus interwork with a border router which is similarly enabled, although in variants these may be omitted. The special IPv6 addresses 94 (see Figure 6) used for this technique have a three-part format of which a first part 96 having thirty two bits is a prefix uniquely identifying that the packet is to be tunnelled by the 6to4 tunnelling technique, a second part 98 having thirty two bits is the IPv4 address of the 6to4 tunnel endpoint, and the third part 100 having sixty four bits is known as the interface ID which is the modified MAC address of the destination host. In variants having a different tunnelling encapsulator, a different respective prefix is used for the same purpose.

[0055] In some variants of these embodiments, the controller 64A is arranged to recognise the presence of this prefix within the retrieved destination address of a received packet and to command the 6to4 tunnelling encapsulator 90A to handle the received packet, and in this case the 6to4 tunnelling encapsulator 90A is arranged to retrieve the special IPv6 address and to extract from its second part 98 the IPv4 address of the 6to4 tunnel endpoint.

[0056] Where, as mentioned above, the controller 64A is arranged to perform prefix recognition, the prefix to be recognised is stored in a storage location of its internal memory 68A, and this storage location can be an entry or part of an entry of the IPv6/tunnel endpoint table 76A.

[0057] The un-encapsulators 88B and 92B have respective IPv4 addresses, which are used by the corresponding encapsulators 86A and 90A in generating their respective encapsulated packets.

[0058] In the preferred arrangement of border router having a plurality of different encapsulators, e.g. 86, 90, the controller 64A accesses the IPv6/tunnel endpoint table 76A in accordance with a set of match criteria to cover the range of possible situations. These are

(a) a tunnel has already been established by the above described DNS Request technique and a IPv6 destination-specific IPv6/IPv4 entry exists in the IPv6/tunnel endpoint table 76A;

(b) a tunnel has already been established by one of the known tunnelling techniques and an IPv6/IPv4 entry exists in the IPv6/tunnel endpoint table 76A, of which the first element 78A has a first part in the form of a specific prefix corresponding to that tunnelling technique;

(c) network management personnel have manually configured the border router 16 to define a tunnel

to a specific IPv6 destination host using 6to4 (or 6over4) to another border router (which might be the border router 16B or a different border router (not shown) associated with a further IPv6 domain (not shown)), and for this case the IPv6/tunnel endpoint table 76A has an entry whose first element 78A is the IPv4-compatible (or IPv4-mapped) address of that destination host;

(d) network management personnel have manually configured the border router 16A to define a tunnel to unspecified IPv6 destination hosts using 6to4 (or 6over4) to another border router (which might be the border router 16B or a different border router associated with a further IPv6 domain), and for this case the IPv6/tunnel endpoint table 76A has an entry having its first element 78A in the form of the 6to4 (or 6over4) prefix followed by the IPv4 address of that other border router followed by null characters, and in some variants the second element 80A of this entry contains null characters, whilst in yet other variants the second element 80A of this entry contains the IPv4 address of that other border router; and

(e) the table contains an entry whose first element 78A is a generic IPv4-compatible or IPv4-mapped IPv6 address, i.e. its first eighty bits are all zeros and the final thirty two (or in a variant, forty eighty) bits are null characters (zeros), the second and fourth elements 80A and 84A contain null characters, and a third element 82A contains the identifier "PC".

[0059] The controller 64A uses its IPv6/tunnel endpoint table 76A to determine the appropriate handling of a received packet in the following manner.

[0060] If the controller 64A finds an entry having a first element 78A matching the complete retrieved destination address, then the content of the second element 80A of that entry is retrieved and used as the IPv4 destination address, i.e. that of the border router 16B, the tunnel endpoint. Additionally, the content of the third element 82A of that entry is retrieved and used as a check that the retrieved IPv4 destination address and the packet received by the border router 16A are to be processed by the encapsulator 86A. The content of the third element 82A is either an identifier for an encapsulator 86A, 90A (e.g. "EN") or an identifier for the protocol converter 32A (e.g. "PC"). As a further check, the fourth element 84A of that entry contains an identifier for the encapsulation type. In other words, for an entry matching the complete retrieved destination address, as just described, the encapsulation type identifier will be "DNS" to signify that the encapsulator 86A is to be used.

[0061] If the controller 64A finds an entry whose first element 78A has the first thirty two bits matching the first thirty two bits, i.e. the special 6to4 prefix part, of the retrieved destination address, then the third and fourth elements 82A and 84A of that entry are checked ("EN"

and "6to4", respectively), the IPv4 destination address is retrieved from the second element 80A of that entry and sent with the packet received by the border router 16A to be processed by the encapsulator 90A.

[0062] If the retrieved destination address is either IPv4-compatible or IPv4-mapped, i.e. the packet is to be protocol converted for an IPv4 destination, then its first eighty bits will be all zeros, and the following sixteen bits will be either all zeros if the address is IPv4-compatible, or all ones if the address is IPv4-mapped. Thus the controller 64A checks to see whether its IPv6/tunnel endpoint table 76A contains an entry whose first element 78A has its first eighty bits all zeros. The content of the second element 80A of such an entry will be all null characters, because no tunnel is involved, and the content of the fourth element 84A of such an entry will be all null characters, because no encapsulation is involved. The content of the third element 82A of that entry is retrieved and used as a check that the retrieved IPv4 destination address and the packet received by the border router 16A are to be processed by the protocol converter 32A. In this case, the content of the third element 82A is an identifier for the protocol converter 32A (e.g. "PC").

[0063] In other variants, the controller 64A is arranged to access the IPv6/tunnel endpoint table 76A, as before, and only command the 6to4 tunnelling encapsulator 90A upon detecting a match with an entry, and in this case the controller 64A passes the special IPv6 address to the 6to4 tunnelling encapsulator 90A, or alternatively the controller 64A extracts the IPv4 address of the 6to4 tunnel endpoint and passes it to the 6to4 tunnelling encapsulator 82A, or yet again the controller 64A, upon detecting such a match, retrieves the element 80A of the matching entry. This element 80A contains the IPv4 address of the 6to4 tunnel endpoint, which was inserted into that element 80A by the controller (or manually) upon creation of that entry.

[0064] In the above embodiment, the local DNS server for the source host 28 is the IPv6 DNS server 24, but in alternative embodiments it can be one of the IPv4 servers 22 of the DNS 20. In such alternative embodiments, although the host 28 can send a DNS Request message for obtaining the IPv6 address of the host 30 and, by the present invention, establish a tunnel across the IPv4 domain 10, the situation is not symmetrical in that the host 30 cannot act as a source and establish a corresponding tunnel across the IPv4 domain 10 to the host 28. For an IPv6 host to be contactable, i.e. act as destination, by the method of the present invention, it is necessary for the DNS server local to that IPv6 host to be an IPv6 DNS server on the same IPv6 domain as that IPv6 host, because the DNS Response message has to pass through the border router adjacent to that IPv6 host in order that the tunnel can be established. In other words, the DNS Request message has to pass through the IPv4 domain and not stop at an IPv4 DNS server acting as the local DNS server for the intended destina-

tion IPv6 host.

Claims

1. A method of establishing a tunnel from a first interface between a first network and a second network to a second interface between the second network and a third network, the first and third networks operating in accordance with a first transmission protocol and having addresses in accordance with a first addressing convention, and the second network operating in accordance with a second transmission protocol and having addresses in accordance with a second addressing convention, the tunnel being for the transport of messages from a first host on the first network to a second host on the third network, the method comprising the steps of:

5 sending from the first host an address request message in accordance with the first transmission protocol, referred to herein as a first type address request message, containing the name of the second host;

10 upon receipt of the address request message at a name to address conversion system of the third network, returning an address response message in accordance with the first transmission protocol, referred to herein as a first type address response message, and containing the address of the second host in a response address field;

15 upon receipt of that first type address response message at the second interface,

20 converting it to an address response message in accordance with the second transmission protocol, referred to herein as a second type address response message, and augmenting that converted second type address response message by fields respectively containing the address of the second interface in accordance with the second addressing convention and the address of the second host in accordance with the first addressing convention; and

25 upon receipt of that augmented that converted second type address response message at the first interface,

30 converting it to a first type address response message, retrieving the contents of the augmenting fields, storing at the first interface a mapping of the retrieved address of the second host

and the retrieved address of the second interface for use in encapsulating messages from the first host addressed to the second host, and

35 replacing the content of the response address field of the resulting first type address response message by the retrieved address of the second host.

2. A method as claimed in claim 1, wherein for establishing a tunnel in the reverse direction for the transport of messages from the second host to the first host, the method comprises the further steps of:

40 upon receipt at the first interface of a message from the first host addressed to the second host, encapsulating that received message in accordance with the first mapping; and upon receipt of the encapsulated message at the second interface,

45 un-encapsulating that received encapsulated message, retrieving from the encapsulating header the address of the first interface in accordance with the second addressing convention, retrieving from the un-encapsulated message the address of the first host in accordance with the first addressing convention, and storing at the second interface a mapping of the retrieved address of the first host and the retrieved address of the first interface for use in encapsulating messages from the second host addressed to the first host.

3. A method as claimed in either claim 1 or claim 2, including setting a time to live for a said stored mapping, and rendering that stored mapping unuseable upon the expiry of the time to live.

4. A method as claimed in claim 3, wherein the rendering step deletes that stored mapping.

5. A method of sending packets from a first host on a first network via a second network to a second host on a third network, the first and third networks operating in accordance with a first transmission protocol and having addresses in accordance with a first addressing convention, and the second network operating in accordance with a second transmission protocol and having addresses in accordance with a second addressing convention, comprising the steps of:

55 establishing a tunnel from a first interface between the first network and the second network

to a second interface between the second network and the third network in accordance with the method of claim 1;

upon receipt by the first host of the resulting first type address response message from the first interface, retrieving the content of the response address field;

generating one or more packets for transmission having a header including source and destination address fields, the source address field containing the address of the first host, and the destination address field containing the retrieved content of the response address field; sending the or each generated packet to the first interface;

for the or each said generated packet received by the first interface, accessing the stored mappings in accordance with the destination address of the received generated packet,

retrieving the stored interface address of the mapping whose retrieved host address matches the destination address of the received generated packet,

generating an encapsulated packet having a payload formed by the received generated packet, and having a header including source and destination address fields, the source address field containing the address of the first interface, and the destination address field containing the retrieved interface network address,

sending the encapsulated packet to the second interface; and

for the or each encapsulated packet received by the second interface, un-encapsulating the received encapsulated packet to recover the original generated packet forming its payload, and

sending that recovered packet to the second host.

6. A method as claimed in claim 5, including the step of storing the retrieved content in association with the name of the second host.

7. A method as claimed in claim 6, including the steps of setting a time to live for the stored retrieved content, and rendering that stored retrieved content un-useable upon the expiry of the time to live.

8. A method as claimed in claim 7, wherein the rendering step deletes the stored retrieved content.

9. A method of establishing a tunnel from a first interface between a first network and a second network to a second interface between the second network

and a third network, the method being substantially as hereinbefore described with reference to the drawings.

10. A method of sending packets from a first host on a first network via a second network to a second host on a third network, the method being substantially as hereinbefore described with reference to the drawings.




Fig.1

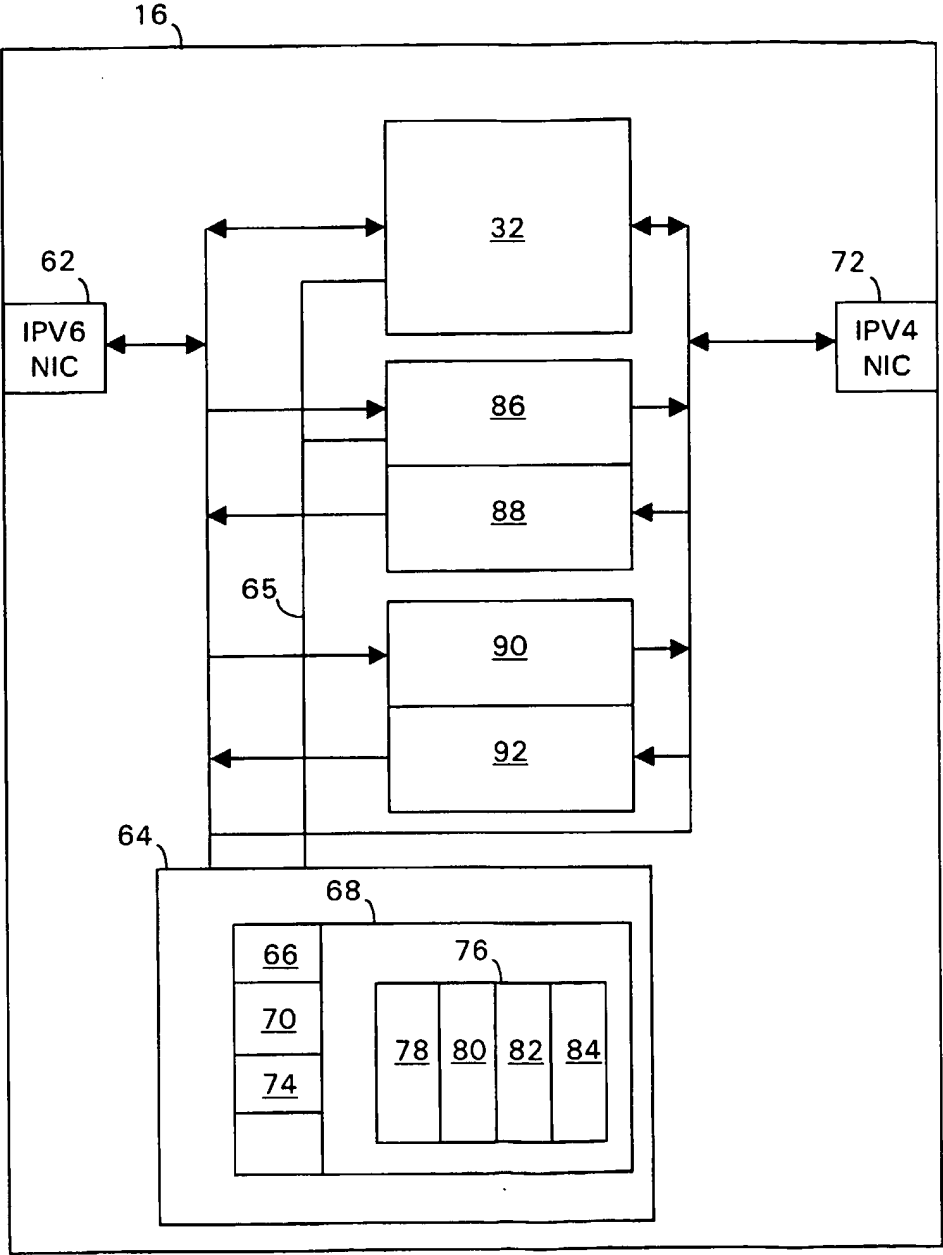


Fig. 2




Fig. 3




Fig. 4




Fig. 5




Fig. 6



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 99 30 7550

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
A	US 5 802 285 A (HIRVINIEMI SEPPÖ) 1 September 1998 (1998-09-01) * abstract * * page 2, right-hand column, line 13 - page 3, left-hand column, line 28 *	1-10	H04L12/46
A	LABORDE D: "UNDERSTANDING AND IMPLEMENTING EFFECTIVE VPNS" COMPUTER TECHNOLOGY REVIEW, US, WESTWORLD PRODUCTION CO. LOS ANGELES, vol. 18, no. 2, 1 February 1998 (1998-02-01), page 12, 14, 16 XP000733729 ISSN: 0278-9647 * the whole document *	1-10	
A	EP 0 840 482 A (HITACHI LTD) 6 May 1998 (1998-05-06) * claims *	1	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
			H04L
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 22 September 2000	Examiner Perez Perez, J
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1503 (03.02) (P04001)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 99 30 7550

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

22-09-2000

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5802285 A	01-09-1998	FI 90710 B GB 2267418 A,B	30-11-1993 01-12-1993
EP 0840482 A	06-05-1998	JP 10136052 A JP 11055319 A	22-05-1998 26-02-1999

EPO FORM P/458

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82